Using Off-the-Shelf Exception Support Components in C++ Verification

Vladimír Štill
Faculty of Informatics,
Masaryk University
Brno, Czech Republic
Email: xstill@fi.muni.cz

Petr Ročkai
Faculty of Informatics,
Masaryk University
Brno, Czech Republic
Email: xrockai@fi.muni.cz

Jiří Barnat
Faculty of Informatics,
Masaryk University
Brno, Czech Republic
Email: barnat@fi.muni.cz

arXiv:1703.02394v1 [cs.SE] 7 Mar 2017

***Abstract*—An important step toward adoption of formal methods in software development is support for mainstream programming languages. Unfortunately, these languages are often rather complex and come with substantial standard libraries. However, by choosing a suitable intermediate language, most of the complexity can be delegated to existing execution-oriented (as opposed to verification-oriented) compiler frontends and standard library implementations. In this paper, we describe how support for C++ exceptions can take advantage of the same principle. Our work is based on DiVM, an LLVM-derived, verification-friendly intermediate language.**

Our implementation consists of 2 parts: an implementation of the `libunwind` platform API which is linked to the program under test and consists of 9 C functions. The other part is a preprocessor for LLVM bitcode which prepares exception-related metadata and replaces associated special-purpose LLVM instructions.

***Index Terms*—Model Checking, Exceptions, C++, Unwinder**

I. Introduction

Today, formal verification methods are not commonly used in software development, even though they are superior to traditional testing approaches in many respects. One particular example is model checking, which can be used to control non-determinism in programs, especially when it arises from parallelism. Formal methods can also be used to extend testing coverage (e.g. via systematic fault injection), verification of liveness properties or verification of global safety properties (such as global assertions). However, to make those advantages actually available to software developers, verification tools must be easy to integrate into existing workflows. If the use of verification tools requires substantial effort (as compared to testing), the costs associated with formal methods can outweigh the savings they provide. This is especially true for modern development processes (especially in commodity software), where there is little time for a separate modeling and design phases.

For this reason, both the academic and industrial communities [2] increasingly seek to develop and use tools which work with mainstream programming languages. However, support for these programming languages – especially when compared to special-purpose modelling formalisms – brings new complexity to verification tools. Programs written in such languages are usually more complex and on a lower level of abstraction than models specifically built for analysis tools. Additionally, many programming languages contain features with no counterparts in a typical modeling language, such as dynamic memory, run-time type information and introspection (RTTI), exception handling, or template instantiation. Moreover, programs written in these languages usually make use of extensive standard libraries. Therefore, the verifier either has to include all of the language and library functionality as primitives, or it has to provide an implementation which is added to the verified program just like a traditional library.

The paper is structured as follows: the remainder of Section I gives motivation, context and contribution of this work. Section II describes the mechanisms that C++ implementations typically use in order to support exceptions. The following Section III then details how LLVM is interpreted in DIVINE 4, in particular the parts relevant to exception handling, such as the stack layout. Section IV and Section V discuss the new components: the LLVM transformation and the unwinder, respectively. Section VI surveys the related work and finally, in Section VII, we evaluate our approach and we summarise our findings in Section VIII.

A. Motivation

In many cases, it is impractical to re-implement the entire programming language and its support libraries. Verification tools can, however, take advantage of existing compilers or libraries to deal with some of the complexity. For example, verification can be substantially simplified by translating the source code into an intermediate representation (IR) using an existing compiler frontend. If the compiler in question can emit intermediate representation after it has been optimised, the verification result is independent of the correctness of the (complex and error-prone) optimiser: any problems introduced by the optimiser will be caught by the verification tool.

In case of C++, a suitable frontend is the clang compiler, which uses LLVM as its IR. Since LLVM can optimise the IR and produce executable code on many platforms from a single optimised IR file, the verification effort does not need to be repeated for each target platform separately. Of course, the code generator (which is comparatively simple when compared

This work has been partially supported by the Czech Science Foundation grant No. 15-08772S and by Red Hat, Inc.

to the platform-neutral optimiser) still needs to preserve the semantics of the program – otherwise, it would invalidate the verification result.

As an alternative to re-using finished, execution-oriented components, one could only support a subset of a programming language (i.e. exclude the parts that are hard to support in a verification tool). However, this weakens the case for supporting mainstream programming languages: it prevents developers from verifying production code. This is especially true for standard libraries, as programming without them requires the programmer to implement everything from scratch. Finally, the standard library is often implemented in the programming language it is part of, and is therefore another good candidate for sharing code with execution-oriented implementations of the language. Unfortunately, upstream implementations of standard libraries usually make extensive use of advanced language features. Consequently, in order to re-use existing standard library implementations, more complete language support is required in the verifier.

Exceptions are among the features that are both widely used (including by the standard library) and tricky to implement. Their use is, however, also common outside of the standard library: libraries like `boost` and application-level code often take advantage of this capability. This is natural, since exceptions simplify error handling and usually require less boilerplate code than any of the alternatives. Furthermore, even though many C++ standard library implementations can be built without exception support[1], this change can significantly affect its behaviour (and as such, validity of the verification result). Finally, error handling paths, including exception propagation, are an important target for analysis by verification tools, as they are both hard to test by more conventional means and likely to contain errors – this naturally arises from the fact that their purpose is to handle unlikely side cases which can be hard to accurately reproduce with testing. A model checker, on the other hand, can take advantage of its built-in support for non-determinism to rigorously explore error paths.[2]

B. Component Re-Use

Unfortunately, off-the-shelf components from execution-oriented language kits do not provide a complete toolbox that would allow verification tool developers to simply concentrate on verification. The difficulties roughly fall into two categories: first, the components interact with each other and with the system for which they were originally designed and second, it is often not at all obvious which components are suitable for re-use and which are not. When a component C is re-used, all the interfaces it uses must be provided as well. There are 3 basic ways in which this can be arranged:

1. re-use another component, D, which provides this interface; this is only possible if all interfaces D uses are already available or can be provided
2. modify component C to avoid its dependency on the interface in question
3. re-implement the interface as a new, possibly tool- or verification-specific component

C. Contribution

The main contribution of this paper is twofold: first, we identify the components that are best re-used and those which are best re-implemented and show that this decision crucially depends on the underlying intermediate language. Second, we provide implementations of the components which cannot be re-used in a form that is easy to integrate into both existing and future verification tools. One of the components works as an LLVM transformation pass, and could be used with any LLVM-based tool. The other component targets the DiVM language [12] specifically, and will therefore only work with tools which understand this language.[3]

The goal of this paper, especially in the context of our previous work on the topic of C++ exceptions in verification [14], is to aid authors of verification tools to minimise costs and effort associated with inclusion of exception support. Depending on the characteristics of the tool, either the approach described in [14] or the one in this paper might be more suitable. Overall, in a verifier which can handle the DiVM language or equivalent, the approach given in this paper is simpler to implement and more robust. A more detailed comparison of the two approaches is given in Section VII-B.

All source code related to this paper, along with more detailed benchmark results and other supplementary material, are available online under a permissive open-source licence.[4]

D. Implementation

Our primary implementation platform is the DIVINE model checker [1]. The C++ support in DIVINE has several components: first, DIVINE uses clang to translate C++ into LLVM IR. As outlined above, the verifier does not need to handle complex syntactical features of C++ this way. A few verification-specific transformations are done on the LLVM IR before it is converted into the DiVM language for execution in DIVINE's Virtual Machine. The VM executes instructions and performs safety checks, such as bound checking. Alone, these components provide basic support for C++. In order to support features such as RTTI and exceptions, it is also necessary to provide a runtime support library and an implementation of the standard library. These libraries in turn rely on a C standard library and on a threading library (`pthreads` on POSIX compatible systems). Those libraries are provided by

[1]There are cases where not using exceptions makes sense: if the end-user code makes no use of them but the standard library is compiled with exception support, the requisite metadata tables only serve to increase the size of the compiled program.

[2]This is a form of fault injection. When using a model checker, it is only necessary to modify the function where the error may arise (e.g. the `malloc` function may be modified to return a `NULL` pointer non-deterministically). The model checker will then take care of exploring all possible combinations of succeeding and failing memory allocations in the program.

[3]DiVM is a relatively small extension of the LLVM IR, therefore extending tools which work with pure LLVM to also support DiVM may be quite easy.

[4]https://divine.fi.muni.cz/2017/exceptions

DiOS, a small, verification-oriented operating system which runs inside DiVM.

As discussed above, building those libraries into the verifier is impractical due to cost and time constraints. There is, however, another important reason why these should be kept out of the verification core: any extension of the verifier increases risks of implementation errors, and the more complex these extensions are, the higher are the associated risks. Moreover, any such errors in the verifier can lead to incorrect verification results. For this reason, DIVINE ships source code implementing these libraries as separate modules; this source code is later compiled into LLVM IR and linked to the verified program. This way, the libraries are subject to the same error checking as user code, and any errors in their implementation that are exposed by the user program will be detected by the verifier.

Additionally, whenever off-the-shelf components are re-used, it is preferable to keep verification-specific changes at minimum. The standard C library in DIVINE is based on PDCLib, a small, portable, public domain C library. The copy of PDCLib in DIVINE includes a few modifications (the C library interfaces directly with the operating system in many cases, therefore it is necessary to port it to work with the verifier, much like it would be necessary to port it to a new operating system). For threading support, DIVINE ships with a custom implementation of the `pthread` library (so far, no existing implementation of the `pthread` interface which could be re-used has been identified). For C++ support, `libc++abi` (the runtime library) and `libc++` (the standard library) are used. Both of these libraries are maintained by the LLVM project and work on many Unix-like systems.

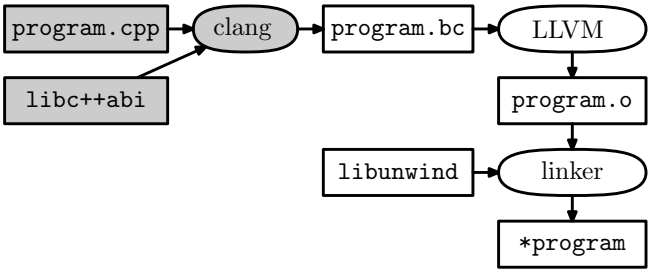

Fig. 1. Components involved in exception support in the standard clang/LLVM stack. Under the scheme proposed in this paper, the highlighted elements are shared between verification and execution environments.

E. Components for Exception Support

Unlike other features of C++, exceptions are neither handled by the standard or runtime libraries alone, nor delegated to the C standard library (as C has no support for exceptions). Instead, `libc++abi` provides exception support with the help of a platform-specific *unwinder library* which is responsible for stack introspection and unwinding (removal of stack frames and transfer of control to exception-handling code). The interaction of these components is illustrated in Figure 1.

For this reason, DIVINE has to either provide an unwinder implementation compatible with `libc++abi`, or modify `libc++abi` to use custom code for exception handling. In DIVINE 3, the latter approach was used, as it was deemed easier at the time [14]. However, while basic exception support was easier to achieve this way, the approach also had its disadvantages. First, the LLVM interpreter in DIVINE 3 had special support for exception-related functionality. Second, the `libc++abi` code for exception handling was replaced, which had 2 important consequences: first, the replacement code was not comprehensive enough[5] and second, this meant that the replaced part of `libc++abi` was not taken into account during verification.

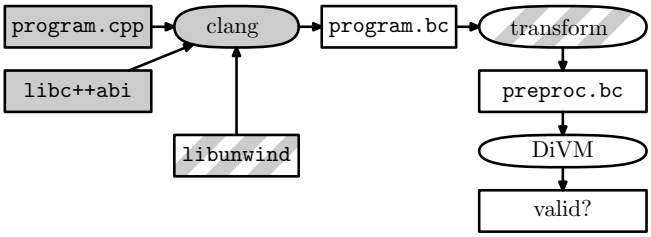

Fig. 2. Components involved in exception support in the DIVINE 4 C++ verification stack. The solid-filled elements are re-used without modification from the execution-oriented clang/LLVM stack (cf. Figure 1). The hatch-filled components are the additions described in this paper.

In this paper, we instead take the first approach: re-use `libc++abi` in its entirety and provide the interfaces it requires. Therefore, we have implemented the `libunwind` interface used by `libc++abi` for stack unwinding and an LLVM transformation which builds metadata tables that `libc++abi` needs to decide which exceptions should be caught, how they should be handled and which functions on the stack need to perform cleanup actions. The situation is illustrated in Figure 2.

Using the original `libc++abi` code means that all features of the C++ exception system are fully supported and verification results also cover the low-level exception support code. That is, this portion of the code is identical in both the bitcode which is verified and in the natively executing program.[6] Finally, the proposed design is easier to extend to other programming languages.

F. Other Components in Use

In line with the principles outlined so far, the implementation of the C and C++ standard libraries (and the C++ runtime library) used in DIVINE are third party code with only minimal modifications. The C standard library implementation (PDCLib) consists of approximately 38 thousand lines of code, while the C++ runtime library (`libc++abi`) and the C++ standard library (`libc++`) contain 8 and 12 thousand lines of code, respectively.

[5]That is, some of the less frequently used features of C++ exceptions were handled either incorrectly or not at all. That is to say, the size of the `libc++abi` portion that would have needed to be re-implemented was initially underestimated.

[6]Clearly, the `libunwind` implementation is different in those two environments, and therefore correctness of the platform-specific implementation of `libunwind` must be established separately.

Standard libraries inevitably contain platform-specific code, and this is also true of the implementations bundled with DIVINE. The modifications due to the porting effort were, however, quite minimal, since DiOS already provides a very POSIX-like interface. The C library was, unsurprisingly, affected the most: changes in memory allocation, program startup- and exit-related functions and in handling of the `errno` variable were required. In `libc++`, however, the changes were limited to platform configuration macros and the only change in `libc++abi` was a DiOS-specific tweak in allocation of thread-local storage for exception handling.

Since user programs and libraries alike rely on the POSIX threading API (also known as `pthread`), this API is provided by DiOS and is implemented in about 2000 lines of C++. The `libunwind` implementation introduced in this paper brings in additional 350 lines of code (the implementation is done in exception-free C++). Likewise, the C library and everything above also depends on low-level filesystem access routines provided by the operating system. In DiOS, this IO and filesystem layer (VFS[7]) is implemented in about 5500 lines of C++ code and uses exceptions heavily for error propagation.

So far, all the components mentioned in this section are linked with the user program to form the final bitcode file for verification. For comparison, the verification core of DIVINE (the DiVM evaluator, memory management and the verification algorithm), amounts to roughly 6 thousand lines of C++. Finally, there is about 2500 lines of code which implement various transformations on the LLVM bitcode. Out of these 2500 lines, less than 300 are part of the exception-related extension described in this paper.

II. Exceptions in C++

Throwing an exception requires removal of all the stack frames[8] between the throwing and catching function from the stack (*unwinding*). Therefore, exception handling is closely tied to the particular platform and is described by ABI[9] for the platform. Commonly, exception handling is split into two parts, one which is tied to the platform (the *unwinder library* which handles stack unwinding) and one which is tied to the language and provided by the language's runtime library.[10] These two parts cooperate in order to provide exception handling for a given language; however, this communication is not standardised in any cross-platform fashion. For this reason, we will now focus on zero-cost exceptions based on the Itanium ABI, an approach which is used across various Unix-like systems on `x86` and `x86_64` processor architectures and is the preferred basis for LLVM exceptions. Nevertheless, it is possible to generalize our results to other implementations.

A. Zero-Cost Exceptions

The so-called zero-cost exceptions are designed to incur no overhead during normal execution, at the expense of relatively costly mechanism for throwing exceptions. This in particular means that no checkpointing is possible. Instead, when an exception is thrown, the exception support library, with the help of *unwind tables*, finds an appropriate *handler* for the exception and uses the *unwinder* to manipulate the stack so that this handler can be executed. The search for the handler is driven by a *personality function*, which is provided by the implementation of the particular programming language.

The personality function is responsible for deciding which handler should execute (the handler selection can be complex and language-specific). In general, there are two types of handlers, *cleanup handlers*, which are used to clean up lexically scoped variables (and call their destructors, as appropriate) and *catch handlers*, which contain dedicated exception-handling code. The latter typically arise from `catch` blocks. Another major difference between those two types of handlers is that catch handlers stop the propagation of the exception, while cleanup handlers let propagation continue after the cleanup is performed. While cleanup handlers are usually run unconditionally, the catch handler to be executed, if any, is determined by the personality function.[11] In C++, the personality function selects the closest `catch` statement which matches the thrown type (the match is determined dynamically, using RTTI). The personality function consults the unwind tables, in particular their *language-specific data area (LSDA)*, to find information about the relevant catch handlers.

When an exception is thrown, the runtime library of the language creates an *exception object* and passes it to the unwinder library. The actual stack unwinding is, on platforms which build on the Itanium ABI, performed in two phases. First, the stack is inspected (without modification) in search for a catch handler. Each stack frame is examined by the relevant personality function.[12] If an appropriate catch handler is found in this phase, unwinding continues with a second phase; otherwise, an unwinder error is reported back to the throwing function. Unwinder errors usually cause program termination. In the second phase, the stack is examined again, and a personality function is invoked again for each frame. In this phase, cleanup handlers come into play. If any handler is found (cleanup or catch), this fact is indicated to the unwinder, which performs the actual unwinding to the flagged frame. Once the control is transferred to the handler, it can either perform cleanup and resume propagation of the exception, or,

[7] Short for Virtual File System, since in a verification environment, the system under test must not access the real filesystem or any other part of the outside environment.

[8] The execution stack of a (C++) program consists of stack frames, each holding context of a single entry into some function. It contains local variables, a return address and register values which need to be restored upon return.

[9] *Application Binary Interface*, a low-level interface between program components on a given platform.

[10] There are many implementations of the C++ runtime library, which, besides exception support code, provides additional features such as RTTI. Each implementation is usually tied to a particular C++ standard library. Commonly used implementations on Unix-like systems are `libsupc++`, which comes with `libstdc++` and the GCC compiler, and `libc++abi`, which is tied to `libc++` used by some builds of clang and by DIVINE.

[11] In fact, the personality function can also decide to skip cleanup handlers, but this is not common.

[12] Different personality functions can be called for different frames, for example if the program consists of code written in different languages with exception support.

if it is a catch handler, end the propagation of the exception. If exception propagation is resumed, the unwinder continues performing phase 2 from the point of the last executed handler. This is facilitated by storing the state of the unwinder within the exception object.

B. Unwind Tables

As mentioned in Section II-A, both the unwinder library and the language runtime depend on unwind tables for their work. The unwinder uses these tables to get information about stack layout in order to be able to unwind frames from it, and for detection which personality function corresponds to a frame. The personality function then uses the language-specific data area (LSDA) of these tables in its decision process.

While the unwinder part of the tables is unwinder- and platform-specific (it depends on stack layout), the LSDA is platform- and language-specific. For these reasons, unwind tables are not present in the LLVM IR; instead, they are generated by the appropriate code generator for any given platform, based on information in the `landingpad` instructions, and the personality attribute of functions. On Unix-like systems, the unwind tables are in the DWARF[13] format.

III. Execution of LLVM programs

In this section, we will look at how LLVM bitcode is executed by a model checker and how this execution is affected by addition of exception support. Unlike previous approaches, the technique described in this paper does not require any exception-specific intrinsic functions or hypercalls to be supported by the verifier. The exception-specific LLVM instructions can be implemented in the simplest possible way: `invoke` becomes equivalent to a `call` instruction followed by an unconditional branch. The `landingpad` instruction can be simply ignored by the verifier and `resume` instructions and calls of the `llvm.eh.typeid.for` intrinsic are both removed by the transformation described in Section IV. Moreover, the metadata required by `libc++abi` are likewise generated by the LLVM transformation and this process is completely transparent to the verifier.

In addition to support for LLVM, the unwinder (described in more detail in Section V) requires the ability to traverse and manipulate the stack and read and write LLVM registers associated with a given stack frame. Finally, it needs access to a representation of the bitcode for a given function. All those abilities are part of the DiVM specification [12] and are generally useful, regardless of their role in exception support.

The DiVM implementation in DIVINE 4 handles execution of LLVM instructions, LLVM intrinsic functions and DiVM-specific *hypercalls*.[14] Hypercalls exist to allocate memory, perform nondeterministic choice or to set DiVM's *control registers* (which contain, among other, the pointer to the currently executing stack frame). Additionally, DiVM performs safety checks, such as memory bound checking, and detects use of uninitialised values. However, DiVM hypercalls are intentionally low-level and simple and do not provide any high-level functionality, such as threading or standard C library functionality. Instead, those are provided by the DIVINE Operating System (DiOS) and the regular C and C++ standard libraries.

The most important purpose of DiOS is to provide threading support. To this end, DiOS provides a *scheduler*, which is responsible for keeping track of threads and their stacks and for (nondeterministically) deciding which thread should execute next. This scheduler is invoked repeatedly by the verifier to construct the state space. The scheduler fully determines the behaviour (or even presence) of concurrency in the verified program: while DiOS provides asynchronous, preemptive parallelism typical of modern operating systems, it is also possible to implement cooperative or synchronous schedulers instead.

A. Stack Layout and Control Registers

A DiVM program can have multiple stacks, but only one of them can be active at any given time (a pointer to the active stack is kept in a DiVM control register). The active stack is normally either the kernel stack or the stack that belongs to the active thread which was selected by the scheduler. Switching of stacks (and program counters) is performed by the `control` hypercall which manipulates DiVM control registers.

Traditionally, stack is represented as a continuous block of memory which contains an activation frame for each function call. In DiVM, the stack is not continuous; instead, it is a singly-linked list of activation frames, each of which points to its caller. This has multiple advantages: first, it is easy to create a stack frame for a function, for example when DiOS needs to create a new thread; additionally, the linked-list-organized stack is a natural match for the graph representation of memory which DiVM mandates, and therefore can be saved more efficiently [12]. Additionally, this way the stack may be nonlinear, and the unwinder can use this feature to safely transfer control to a cleanup block while the unwinder frame is still on the stack. Later, the handler can return control to the unwinder frame and the unwinder can continue its execution. This would be impossible with a continuous stack since cleanup code is allowed to call arbitrary functions and frames of those functions would overwrite the frame of the unwinder. For this reason, on traditional platforms, the unwinder needs to store its entire state in the exception object, while in DiVM, it can simply retain its own activation frame. An illustration of how the stack looks while the unwinder is active is shown in Figure 3.

IV. The LLVM Transformation

The C++ runtime library (`libc++abi` in our case), needs access to the LSDA section of unwind tables (a pointer

[13] DWARF is a standard for debugging information designed for use with ELF executables. It is used on most modern Unix-like systems.

[14] Intrinsic functions are provided by LLVM as a light-weight alternative to new instructions. Such functions are recognized and translated by LLVM itself, as opposed to "normal" functions that come from libraries or the program. Likewise, DiVM provides hypercalls, which are functions that are, in addition to LLVM intrinsics, recognized by DiVM.

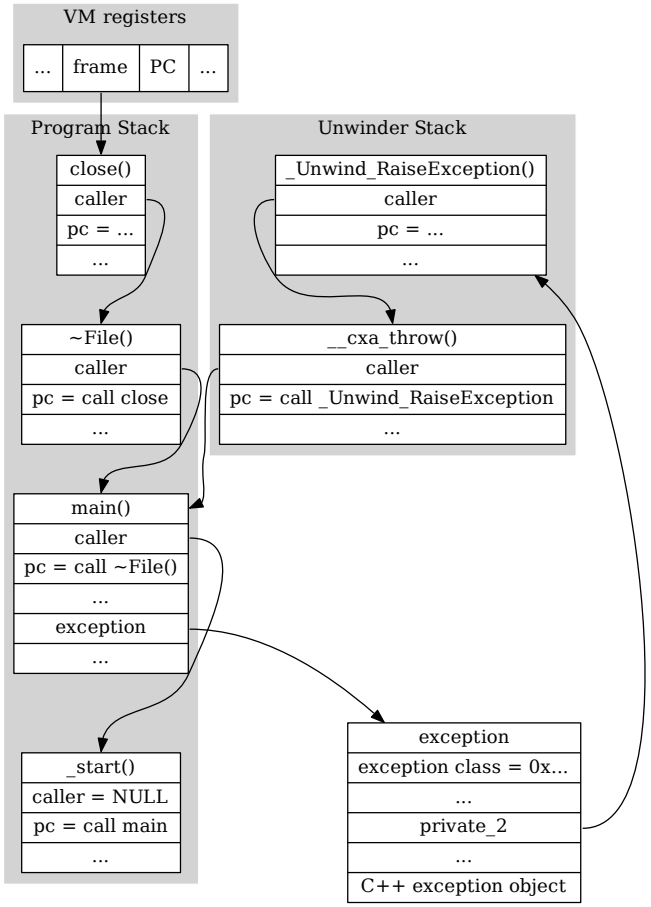

Fig. 3. In this figure we can see a stack of a program which is running cleanup block in the `main` function. The cleanup block calls the destructor of `File` structure, which in turn calls the `close` function (which is the current active function). Furthermore, the cleanup handler can access the exception object which contains a pointer to the stack of the unwinder. This pointer is used by the implementation of the `resume` instruction to jump back to the unwinder and continue phase 2 of the unwinding.

to this metadata section is accessible through the unwinder interface). This section contains DWARF-encoded exception tables, which are normally generated together with the executable by the compiler backend (code generator). Unfortunately, the generator of DWARF exception tables in LLVM is closely tied to the machine code generator and cannot be used to generate DWARF-formatted exception tables for verification purposes. For this reason, we have implemented a small LLVM transformation which processes the information in `landingpad` instructions and generates LLVM constants which contain the DWARF-formatted LSDA data. A reference to one such constant is attached to each function in the bitcode file.

To improve efficiency, LLVM does not directly use RTTI type info pointers within the landing blocks to decide which exception handlers should run. RTTI objects are special C++ objects which are used to identify types at runtime and are emitted by the C++ frontend as constants. Due to the complexities of C++ type system, matching RTTI types against

Function	Description
`SetGR`	Store a value into a general-purpose register
`GetGR`	Read a value from a general-purpose register
`SetIP`	Stora a value into the program counter
`GetIP`	Read the value of the program counter
`RaiseException`	Unwind the stack
`Resume`	Continue unwinding the stack after a cleanup
`DeleteException`	Delete an exception object
`GetLSDA`	Obtain a pointer to the LSDA
`GetRegionStart`	Obtain a base for relative code pointers

TABLE I
A LIST OF C FUNCTIONS PROVIDED BY `LIBUNWIND`. IN C, ALL THE FUNCTIONS ARE PREFIXED WITH `_UNWIND_` TO PREVENT NAME CONFLICTS WITH USER CODE AND OTHER LIBRARIES (I.E. THE C NAME OF `SETGR` IS `_UNWIND_SETGR`).

each other is expensive: a search in a pair of directed acyclic graphs is required. Moreover, since the RTTI matching must be already done in the personality function to decide which frames to unwind, the personality can also pre-compute a numerical index for the landing pad. This index, also called a *selector value* is then used as a shortcut to run an appropriate `catch` clause within the landing block, instead of re-doing the expensive RTTI matching. Since the `catch` handler is typically expressed in terms of typeinfo pointers, it needs to efficiently obtain the selector value from a type info pointer. For this purpose, LLVM provides a `llvm.eh.typeid.for` intrinsic, which obtains (preferably at compile time) the selector value corresponding to a particular type info pointer.

Therefore, besides generating the LSDA data, the transformation statically computes the values which correspond to `llvm.eh.typeid.for` calls and substitutes them into the bitcode. Since the purpose of `llvm.eh.typeid.for` is to translate from RTTI pointers to selector values, it is only required that the integer selector value chosen for a particular RTTI object is in agreement with the personality function. In our implementation, this is ensured by computing the selector values statically for both the LSDA (which is where personality function obtains them) and for `llvm.eh.typeid.for` at the same time.

Finally, the transformation rewrites all uses of the `resume` instruction to ordinary calls to `Resume`, a function which is part of `libunwind` (see also Table I).

V. THE UNWINDER

The unwinder in DIVINE is designed around the interface described in the Itanium C++ ABI documentation,[15] adopted by multiple vendors and across multiple architectures. The implementation is part of the runtime libraries shipped with DIVINE.[16] The unwinder builds upon a lower-level stack access API which is provided by DiOS under `sys/stack.h`.

Due to the stack layout used in DiVM (a linked list of frames, see also Section III-A), our unwinder is much simpler than usual. The main task of unwinding is handled by the

[15] https://mentorembedded.github.io/cxx-abi/abi.html
[16] `runtime/libc/functions/unwind.cpp`

`RaiseException` function, which is called by the language runtime when an exception is thrown. This function performs the two phase handler lookup described in Section II-A and it adheres to the Itanium ABI specification, with the following exceptions:

i. it checks that an exception is not propagated out of a function which has the nounwind attribute set, and reports verification error if this is the case;
ii. if the exception is a C++ exception and there is no handler for this exception type, the unwinder chooses nondeterministically whether it should or should not unwind the stack and invoke cleanup handlers.

The purpose of the first deviation is to check consistency of exception annotations (arising, for example, from a `nothrow` function attribute as available in GCC and in clang). The second modification allows DIVINE to check both allowed behaviours of uncaught exceptions in C++: the C++ standard specifies that it is implementation-defined whether the stack is unwound (and destructors invoked) when an exception is not caught.[17] Since the program may contain errors which manifest only under one of these behaviours, it is useful to be able to test both of them.

A. Low-Level Unwinding

The primary function of the unwinder described above is to find exception handlers; for the actual unwinding of frames, it uses a lower-level interface provided by DiOS. This interface consists of two functions: `__dios_jump`, which performs a non-local jump, possibly affecting both the program counter and the active frame, and `__dios_unwind`, which removes stack frames from a given stack. `__dios_unwind` is designed in such a way that it can unwind any stack, not only the one it is running on, and is not limited to the topmost frames (effectively, it removes frames from the stack's singly-linked list, freeing all the memory allocated for local variables that belong to the unlinked frames, along with the frames themselves[18]). The unwinder identifies values as local variables by looking at the instructions of the active function – the results of `alloca` instructions are exactly the addresses of local variables.

B. Unwinder Registers

When an exception is propagating, a personality function has to be able to communicate with the code which handles the exception. In C++, the communicated information includes the address of the exception object and a selector value which is later used by the handler. On most platforms, these values are passed to the handler using registers, which are manipulated using unwinder's `SetGR` function. This function can either set the register directly (if it is guaranteed not to be overwritten before the control is transferred to the handler), or save the value in a platform-specific way and make sure it is restored before the handler is invoked.

[17] Section 15.5, paragraph 9 of the C++ standard [5]

[18] When a function returns normally (due to a `ret` instruction), DiVM takes care of freeing the frame and its local variables (`alloca` memory).

In LLVM (and hence in DiVM), there is no suitable counterpart to the general purpose registers of a CPU; instead, the values set by the personality function should be made available to the program in the return value of the `landingpad` instruction. This, however, requires the knowledge of the expected semantics of these registers. Currently, all users of the unwinder are expected to use the same registers as the C++ frontend in clang. That is, register 0 corresponds to the exception object and register 1 corresponds to a type index. This also directly maps to the return type of `landingpad` instructions and therefore the register values can be saved directly into the LLVM register corresponding to the particular `landingpad` that is about to be executed.

Registers other than 0 and 1 are currently not supported. In LLVM, in line with the above observation about clang and C++, there is a convention that `SetGR` indices correspond to indices into the result tuple of a `landingpad` instruction. As long as this convention is preserved by a particular language frontend and its corresponding runtime library (personality function), it is very easy to extend our unwinder to support this language. Finally, if a language frontend were instead to emit calls to `GetGR` in the handler, registers of this type can be stored in the unwinder `Context` directly.

C. Atomicity of the Unwinder

The unwinder performs rather complex operations and therefore throwing an exception can create many states, even when τ reduction [13] is enabled. However, many of these states are not interesting from the point of view of verification, as the operations performed by the unwinder are mostly thread-local and only the exception handlers (and possibly personality function) can perform globally visible actions. For this reason, the unwinder uses DiVM's atomic sections to hide most of its complexity.

Since an atomic section is implemented as an *interrupt mask* (i.e. a single flag indicating that an atomic section is executing) in DiVM, it is necessary to correctly maintain the state of this flag. In particular, it is required that the unwinder behaves reasonably even if it is called when the program is already in an atomic section. Consequently, care must be taken to restore the state of the atomic mask when the unwinder transfers control to a personality function or an exception handler. When the unwinder is first called, it enters an atomic section and saves the previous value of the interrupt mask. This will be the value the flag will be restored to when a personality function is first invoked. The mask is later re-acquired after the personality function returns and it is restored once more when the first handler is invoked. When the exception handler resumes (using the `resume` instruction), the atomic section is re-entered and its state saved so its state before the resume can be restored again for the next call to a personality function. This way, it is possible to safely throw an exception out of an atomic section, provided that the atomic section is exception-safe (that is, it has an exception handler which ends the atomic section if an exception is propagated out of it).

D. `longjmp` Support

Using the low-level unwinder interface described in Section V-A, it is easy to implement other mechanisms for non-local transfer of control. The functions `longjmp` and `setjmp`, specified as part of C89, are one such example.[19] The `setjmp` function can be used to save part of the state of the program, so that a later call to `longjmp` can restore the stack to the state it was in when `setjmp` was called. This way, `longjmp` can be used to remove multiple frames from the stack. When `longjmp` is called, the program behaves as if `setjmp` returned again, only this time it returns a different value (provided as an argument to `longjmp`).

The DIVINE implementation of `setjmp` saves the program counter and the frame pointer of the caller of `setjmp`. The `longjmp` function then uses this saved state, along with access to the text of the program, to set the return value of the `call` instruction corresponding to the `setjmp`. Afterwards, it unwinds the stack using the low-level stack access API from `sys/stack.h` and transfers control to the instruction right after the call to `setjmp`.

VI. Related Work

Primarily, we have looked at existing tools which support verification of C++ programs. Existence of an implementation is, to a certain degree, an indication that a given approach is viable in practice. We have, however, also looked at approaches proposed in the literature which have no implementations (or only a prototype) available.

A number of verification tools are based on LLVM and therefore have some support for C++. LLBMC [15] and NBIS [6] are LLVM-based bounded model checkers which target mainly C and have no support for exceptions or the C++ standard library. VVT [7] is a successor of NBIS which uses either IC3 or bounded model checking and has limited C++ support, but it does not support exceptions. Furthermore, KLEE [3] and KLOVER [9] are LLVM-based tools for test generation and symbolic execution. KLOVER targets C++ and according to [9] has exception support, but it is not publicly available. On the other hand, KLEE focuses primarily on C and its C++ support is rather limited and it has no exception support.

Both CBMC [4, 8] and ESBMC [11] bounded model checkers support C++ (but neither appears to support the standard library) and they include support for exceptions. However, in CBMC, the support for exceptions is limited to throwing and catching fundamental types.[20] In our survey of tools for verification of C++ programs, ESBMC has by far the best exception support: the latest version can deal with most, but not all[21], types of exception handlers and even with exception specifications. Finally, DIVINE 3 [14] also comes close to full support for exceptions, but lacks support for exception specifications. Overall, this survey suggests that all current implementations of C++ exceptions in verification tools are incomplete and confirms that using an existing, standards-compliant implementation in a verification tool is indeed quite desirable.

Finally, it is also possible to transform a C++ program with exceptions into an equivalent program which only uses more traditional control flow constructs. This approach was taken in [10], with the goal of re-using existing analysis tools without exception support. While this approach is applicable to a wide array of verification tools, it is also incompatible with re-use of existing exception-related runtime library code. As such, it offers a very different set of tradeoffs than our current approach. Moreover, the translation cost is far from negligible, and also affects code that does not directly deal with exceptions (i.e. it violates the zero-cost principle of modern exception handling). Unfortunately, we were unable to evaluate this approach, since there are no publicly available tools which would implement it.

VII. Evaluation

In order to asses the viability of our approach, we have executed a set of benchmarks in various configurations of DIVINE 4. The benchmarks were executed on quad-core Xeon 5130 clocked at 2 GHz and with 16GB of RAM. We have measured the wall time, making all 4 cores available to the verifier.

A. Benchmark Models

The set of models we have used for this comparison consists of 831 model instances, out of which we picked the 794 that do not contain errors. The reason for this is that the execution time is much more variable when a given program contains an error, since the model checking algorithm works on the fly, stopping as soon as the error is discovered.

Majority of the valid models (777) are C++ programs of varying complexity, while the 17 models in the svc-pthread category are concurrent programs written in plain C with pthreads. Since our implementation of the pthread API is done in C++, the impact of exception support on verification of C programs is also relevant. The "alg" category includes sequential algorithmic and data structure benchmarks, the "pv264" category contains unit tests for student assignments in a C++ course, the "iv112" category contains unit tests for concurrent data structures and other parallel programs (again assignment problems in a C++ course), "libcxx" contains a selection of the `libc++` testsuite (with focus on exception support coverage), "bricks" contains unit tests for various C++ helper classes,

[19] Implemented in `runtime/libc/includes/setjmp.h` and `runtime/libc/functions/setjmp/`.

[20] A simple test which throws and tries to catch an exception object crashes CBCM 5.6.

[21] ESBMC 3.0 is unable to determine that an exception ought to be caught when the `catch` clause specifies a type which is a virtual base class in a diamond-shaped hierarchy and the object thrown is of the most-derived type of the diamond. This suggests that ESBMC uses its own implementation of RTTI support code, which is somewhat incomplete, compared to production implementations.

category	#mod	time (D4)	time (D3)	states (D4)	states (D3)
alg	9	3:52	3:51	543.3 k	543.3 k
pv264	13	1:34	1:32	183.0 k	183.0 k
iv112	11	25:58	25:57	3743 k	3743 k
libcxx	425	42:15	42:09	2182 k	2182 k
bricks	292	3:04:25	2:56:55	6271 k	6251 k
divine	3	6:20	6:18	1040 k	1040 k
cryptopals	3	0:01	0:01	1943	1943
llvm	12	36:36	36:27	3865 k	3865 k
svc-pthread	17	16:47	16:41	1685 k	1685 k
total	794	5:21:44	5:13:49	20.1 M	20.0 M

TABLE II
COMPARISON OF THE NEW EXCEPTION CODE WITH A DIVINE-3-STYLE VERSION.

category	#mod	time (D4)	time (stub)	states (D4)
alg	9	3:52	3:52	543.3 k
pv264	13	1:34	1:34	183.0 k
iv112	11	25:58	26:00	3743 k
libcxx	392	41:56	41:54	2179 k
bricks	192	35:30	35:21	2378 k
divine	3	6:20	6:19	1040 k
cryptopals	3	0:01	0:01	1943
llvm	12	36:36	36:28	3865 k
svc-pthread	17	16:47	16:43	1685 k
total	661	2:52:30	2:52:08	16.2 M

TABLE III
COMPARISON OF THE NEW EXCEPTION CODE AGAINST STUBBED EXCEPTIONS. COMPARED TO TABLE II, IN THIS CASE 133 MODELS FAILED DUE TO THE STUBS. STATE COUNTS ARE IDENTICAL FOR ALL MODELS.

including concurrent data structures, "divine" contains unit tests for a concurrent hashset implementation used in DIVINE, "cryptopals" contains solutions of the cryptopals problem set[22], the "llvm" category contains programs from the LLVM test-suite[23] and finally, the "svc-pthread" category includes pthread-based C programs from the SV-COMP benchmark set. In most of the programs, it was assumed that `malloc` and `new` never fail, with the notable exception of part of the "bricks" category unit tests. The tests where `new` failures are allowed are especially suitable for evaluating exception code, in particular where multiple concurrent threads are running at the time of the possible failure.

B. Comparison to Builtin Exception Support

In addition to the approach presented in this paper, we have implemented the approach described in [14] in the context of DIVINE 4. This allowed us to directly measure the penalty associated with the present approach, which is more thorough and less labour-intensive at the same time. Our expectation was that this would translate to slower verification, since the off-the-shelf code is more complex than the corresponding hand-tailored version used in [14]. In line with this expectation, we set the criterion of viability: we would consider a slowdown of at most 10 % to be an acceptable price for the improved verification fidelity, and convenience of implementation. Since other resource consumption (especially memory) of verification is typically proportional to state space size, we have used the number of states explored as an additional metric. The expected effect on the shape (and, by extension, size) of the state space should be smaller than the effect on computation time (most of the additional complexity is related to computing a single transition). We believe that an acceptable penalty in this metric would be about 2 % increase.

As can be seen in Table II, the time penalty on our chosen model set is very acceptable – just shy of 2.6 % – and the state space size is within 1 % of the older approach [14]. We believe that this small penalty is well justified by the superior verification properties of the new approach.

[22]http://cryptopals.com

[23]http://llvm.org/svn/llvm-project/test-suite/trunk/SingleSource/Benchmarks/Shootout

C. Comparison to Stub Exceptions

The second alternative approach is to consider any thrown exception an error, regardless of whether it is caught or not. This can be achieved much more easily than real support for exceptions, since we can simply replace the entire `libunwind` interface with stubs which raise an error and refuse to continue. This approach only works for models which do not actually throw any exceptions during their execution. The results of this comparison are shown in Table III – the verification time is nearly identical and the state spaces are entirely so. This is in line with expectations: in those models, catch blocks are present but never executed. Since the proposed approach does not incur any overhead until an exception is actually thrown, we would not expect a substantial time difference.

D. Comparison to No Exceptions

Finally, the last alternative is to disable exception support in the C++ frontend entirely. In `clang`, this is achieved by compiling the source code with the `-fno-exceptions` flag. In this case, the LLVM bitcode contains no exception-related artefacts at all, but many programs fail to build. Additionally, a number of programs in the "bricks" category contain exception handlers for memory allocation errors[24] and therefore exit cleanly upon memory exhaustion. Even though some of those programs can be compiled with `-fno-exceptions`, they now contain an error (a null pointer dereference) which is not present when they are compiled the standard way. Those programs were therefore excluded from the comparison. The summary of this comparison can be found in Table IV – the time saved for models where `-fno-exceptions` is applicable is again quite small, less than 13 %. In this case, the difference is due to the changes in control flow of the resulting LLVM bitcode. Since `call` is not a terminator instruction

[24]In this case, the handler is installed using `std::set_terminate`, which is available even when `-fno-exceptions` is given. The situation would be similar if only parts of the program were compiled with `-fno-execptions`. In particular, the problem is that the standard library, if compiled with `-fno-exceptions`, cannot throw, and must therefore behave differently in those scenarios, affecting the behaviour of the user program.

category	#mod	time (D4)	time (nxc)	states
alg	1	0:24	0:23	34.2 k
pv264	1	0:00	0:00	57
iv112	10	23:58	22:06	3571 k
libcxx	393	41:57	40:44	2180 k
svc-pthread	17	16:47	15:42	1685 k
total	423	1:23:33	1:19:21	7504 k

TABLE IV
COMPARISON OF THE NEW EXCEPTION SUPPORT AGAINST A CASE WHERE `-fno-exceptions` WAS USED TO COMPILE THE SOURCES AND LIBRARIES. IN THIS CASE, IT WAS ONLY POSSIBLE TO VERIFY 423 MODELS FROM THE SET (I.E. 371 MODELS ARE MISSING FROM THE COMPARISON). STATE COUNTS ARE IDENTICAL FOR ALL MODELS.

(unlike `invoke`), the *local* control flow in a function is negatively affected by the presence of `invoke` instructions: more branching is required, and this slows down the evaluator in DiVM. While it is easy to see if a given program can be compiled with `-fno-exceptions`, it is typically much harder to ensure that its behaviour will be unchanged. For this reason, we do not consider the time penalty in verification of this type of programs a problem.

E. Re-usability

As outlined in Section I-F, the two components directly involved in exception support are comparatively small and well isolated. The LLVM transformation is fully re-usable with any LLVM-based tool. The unwinder, on the other hand, relies on the capabilities of DiVM. However, there is no need for hypercalls specific to exception handling and therefore, the implementation work is essentially transparent to DiVM. The capabilities of DiVM required by the unwinder are limited to the following: linked-list stack representation, runtime access to the program bitcode and 2 hypercalls: `__vm_control` and `__vm_obj_free`. More details about DiVM can be found in [12].

Finally, adding support for a new type of exceptions is also much simpler in this approach – no modifications to DiVM (or any other host tool) are required: only the two components described in this paper may need to be modified.

VIII. CONCLUSION

In this paper, we have discussed an approach to extending an LLVM-based model checker with C++ exception support. We have found that re-using an existing implementation of the runtime support library is a viable approach to obtain complete, standards-compliant exception support. A precondition of this approach is that the verification tool is flexible enough to make stack unwinding possible. The DiVM language, on which the DIVINE model checker is based, has proven to be a good match for this approach, due to its simple and explicit stack representation, along with a suitable set of control flow primitives.

We also performed a survey of tools based on partial or complete reimplementations of C++ exception support routines and found that in each tool, at least one edge case is not well supported. In contrast to this finding, with our approach, all those edge cases are covered "for free", that is, by the virtue of re-using an existing, complete implementation. Contrary to the prediction made in [14], we have found that with a suitable target language, implementing a new unwinder can be relatively simple. The unwinder implementation described in this paper is only about 350 lines of C++ code, while it would be impossible to implement without verifier modifications in DIVINE 3. Therefore, we can conclude that with the advent of the DiVM specification [12] and its implementation in DIVINE 4, re-implementing the `libunwind` API and re-using `libc++abi` became a viable strategy to provide exception support.

REFERENCES

[1] Jiří Barnat, Luboš Brim, Vojtěch Havel, Jan Havlíček, Jan Kriho, Milan Lenčo, Petr Ročkai, Vladimír Štill, and Jiří Weiser. DiVinE 3.0 – an explicit-state model checker for multithreaded C & C++ programs. In *CAV*, volume 8044 of *LNCS*, pages 863–868. Springer, 2013.

[2] Dirk Beyer. Reliable and Reproducible Competition Results with BenchExec and Witnesses Report on SV-COMP 2016. In *TACAS*, pages 887–904. Springer, 2016. ISBN 978-3-662-49673-2. doi: 10.1007/978-3-662-49674-9_55.

[3] Cristian Cadar, Daniel Dunbar, and Dawson R. Engler. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *8th USENIX Symposium on Operating Systems Design and Implementation, (OSDI 2008)*, pages 209–224. USENIX Association, 2008.

[4] Edmund Clarke, Daniel Kroening, and Flavio Lerda. A Tool for Checking ANSI-C Programs. In *TACAS*, pages 168–176. Springer, 2004. ISBN 978-3-540-24730-2. doi: 10.1007/978-3-540-24730-2_15.

[5] ISO C++ Standards Committee. Standard for Programming Language C++. Working Draft N4296. Technical report, ISO IEC JTC1/SC22/WG21, 2014.

[6] Henning Günther and Georg Weissenbacher. Incremental Bounded Software Model Checking. In *SPIN*. ACM, 2014.

[7] Henning Günther, Alfons Laarman, and Georg Weissenbacher. Vienna verification tool: IC3 for parallel software - (competition contribution). In *TACAS*, pages 954–957, 2016. doi: 10.1007/978-3-662-49674-9_69.

[8] Daniel Kroening and Michael Tautschnig. CBMC – C bounded model checker. In *TACAS*, pages 389–391. Springer, 2014. ISBN 978-3-642-54862-8. doi: 10.1007/978-3-642-54862-8_26.

[9] Guodong Li, Indradeep Ghosh, and Sreeranga P. Rajan. KLOVER: A Symbolic Execution and Automatic Test Generation Tool for C++ Programs. In *CAV*, volume 6806 of *LNCS*, pages 609–615. Springer, 2011. ISBN 978-3-642-22109-5.

[10] Prakash Prabhu, Naoto Maeda, Gogul Balakrishnan, Franjo Ivančić, and Aarti Gupta. Interprocedural excep-

tion analysis for C++. In *ECOOP*, volume 6813 of *LNCS*, pages 583–608. Springer, 2011. ISBN 978-3-642-22654-0.

[11] Mikhail Ramalho, Mauro Freitas, Felipe Sousa, Hendrio Marques, Lucas Cordeiro, and Bernd Fischer. SMT-Based Bounded Model Checking of C++ Programs. In *ECBS*, pages 147–156. IEEE Computer Society, 2013. ISBN 978-0-7695-4991-0.

[12] Petr Ročkai and Jiří Barnat. DiVM: Model checking with LLVM and graph memory. 2017. Preliminary version.

[13] Petr Ročkai, Jiří Barnat, and Luboš Brim. Improved state space reductions for LTL model checking of C & C++ programs. In *NASA Formal Methods*, volume 7871 of *LNCS*, pages 1–15. Springer, 2013.

[14] Petr Ročkai, Jiří Barnat, and Luboš Brim. Model checking C++ programs with exceptions. *Science of Computer Programming*, 128:68 – 85, 2016.

[15] Carsten Sinz, Florian Merz, and Stephan Falke. LLBMC: A bounded model checker for LLVM's intermediate representation. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 7214 of *LNCS*, pages 542–544. Springer, 2012. ISBN 978-3-642-28755-8. doi: 10.1007/978-3-642-28756-5_44.